An Adaptive CUSUM Chart for Drift Detection

(Running title: CUSUM Chart for Drift Detection)

Fan Yi and Peihua Qiu*

Department of Biostatistics, University of Florida, Gainesville, FL 32610, USA.

*Corresponding Author, E-mail: pqiu@ufl.edu

Abstract

In practice, sequential processes often have gradual changes in their process distributions over time. This is related to the drift detection problem in statistical process control. In the literature, there have been some existing discussions on this problem. But, most existing methods are designed based on the assumption that the related drift is linear or have another specific pattern. In reality, however, such specified patterns may not be valid. In this paper, we suggest an adaptive CUSUM chart to handle the drift detection problem with a flexible drift pattern. The new method integrates the general framework to construct a CUSUM chart based on the generalized likelihood ratio statistic and estimation of a shift size by the exponentially weighted least square regression procedure. Simulation studies show that the proposed method is effective in various cases considered. The new method is also illustrated using an example about the exchange rates between Indian Rupees and US Dollars. [1]

Key Words: Adaptive CUSUM chart; Drift detection; Exponentially weighted least square; Likelihood ratio; Linear drift; Statistical process control.

1 Introduction

In environmental science, scientists found that the gradual loss of water resource in certain regions (e.g., the Salton Sea region in California) was damaging the local environment and ecosystems for human beings, animals and plants (Shuford et al. 2002, 2004, Tiffany et al. 2007). In financial

[1] The data used in the paper are available from the authors upon reasonable request.

markets, exchange rates between two currencies would change gradually after the gradual adjustment of certain financial policies like interest rates (Zeileis et al. 2010). In such applications, the distribution of the related response variable would have a drift after a certain time point, and it is important to detect such a drift as soon as possible so that some interventions can be implemented in a timely manner. This paper suggests a new and effective method for drift detection.

In the SPC literature, most control charts are developed for detecting shifts in the distribution of a sequential process (Qiu 2014). There are also some existing discussions on detecting a mean drift for a sequential process (e.g., Gan 1992, Shu et al. 2008, Su et al. 2011, Zhou et al. 2010, Zou et al. 2009). For instance, Gan (1992) discussed the performance of the conventional CUSUM chart designed for detecting mean shifts when the process under monitoring actually had a linear mean drift. Shu et al. (2008) suggested a weighted CUSUM chart for detecting patterned mean shifts, including a linear drift. Su et al. (2011) discussed detection of a linear mean drift by using an adaptive EWMA chart. Zhou et al. (2010) suggested a control chart for monitoring a process with patterned mean or variance changes based on the generalized likelihood ratio statistic. Zou et al. (2009) compared several different control charts for detecting linear drifts, including the conventional CUSUM and EWMA charts, the generalized EWMA chart by Han and Tsung (2004), and two control charts based on the generalized likelihood ratio statistic under the linear drift and step shift alternatives. While these research effort is important, the existing methods described above for drift detection still have much room for improvement because of the complexity of the drift detection problem due to the facts that drifts can have many different patterns in practice, the existing methods were developed for detecting one or more specific patterns, and their specified drift patterns may not describe real drifts well in practice.

In this paper, we propose a new method for drift detection. The new method uses the general framework to construct a CUSUM chart based on the generalized likelihood ratio statistic, after the drift size at the current observation time is estimated by an exponentially weighted least square regression procedure. Because the weights in the weighted least square regression procedure decay exponentially fast for past observations, the proposed CUSUM chart is robust to the actual drift pattern. Numerical studies show that this method performs favorably when compared with some representative existing methods.

The remainder of the paper is organized as follows. In Section 2, the proposed new method is described in details. Its numerical performance is evaluated in Section 3 by some simulation examples. Then, the method is demonstrated using an example about the exchange rates between Indian Rupees and US Dollars in Section 4. Several remarks conclude the paper in Section 5.

2 Proposed CUSUM Chart

Let the process observation at time n be denoted as X_n, for $n \geq 1$. When the process is in-control (IC) at n, assume that $X_n \sim N(\mu_0, \sigma_0^2)$, where μ_0 and σ_0^2 are the IC mean and variance. If the process is out of control (OC) at time n, then it is assumed that $X_n \sim N(\mu_0 + \delta_n \sigma_0, \sigma_0^2)$, where $\delta_n \sigma_0$ is the mean shift size. Let $e_n = (X_n - \mu_0)/\sigma_0$ be the standardized value of X_n, for $n \geq 1$. Then, its IC and OC pdf functions are respectively

$$f_0(e_n) = \frac{1}{\sqrt{2\pi}} \exp(-e_n^2/2), \qquad f_1(e_n) = \frac{1}{\sqrt{2\pi}} \exp[-(e_n - \delta_n)^2/2],$$

and the corresponding log likelihood ratio is

$$\log(f_1(e_n)/f_0(e_n)) = \delta_n(e_n - \delta_n/2).$$

To detect such a mean shift, the conventional CUSUM charting statistic (cf., Qiu 2014, Section 4.2.4) is

$$C_n = \max\{0, C_{n-1} + \delta_n(e_n - \delta_n/2)\}, \text{ for } n \geq 1,$$

where $C_0 = 0$. From the above expression, it can be seen that the charting statistic C_n depends on δ_n which is often unknown in practice. To overcome this difficulty, we suggest estimating δ_n by the following exponentially weighted least square regression procedure:

$$\min_{\alpha, \beta \in R} \sum_{i=1}^{n} (e_i - \alpha - \beta i)^2 (1-\lambda)^{n-i}, \tag{1}$$

where $\lambda \in [0,1)$ is a weighting parameter. In (1), a line has been fitted to $\{e_i, 0 \leq i \leq n\}$ by a weighted least square regression procedure, and the weight $(1-\lambda)^{n-i}$ at time i exponentially decays when i moves away from n. So, only observations whose observation times are close to n would receive relatively large weights in this procedure.

Let the solutions to α and β in (1) be denoted respectively as $\widehat{\alpha}$ and $\widehat{\beta}$. Then, a reasonable estimate of δ_n is

$$\widehat{\delta}_n = \widehat{\alpha} + \widehat{\beta} n.$$

After δ_n is replaced by $\widehat{\delta}_n$ in the expression of C_n, our suggested adaptive CUSUM chart for detecting a mean drift is

$$\widehat{C}_n = \max\left[0, \widehat{C}_{n-1} + \widehat{\delta}_n(e_n - \widehat{\delta}_n/2)\right], \text{ for } n \geq 1, \tag{2}$$

where $\widehat{C}_0 = 0$. The chart gives a signal if

$$\widehat{C}_n > h, \tag{3}$$

where $h > 0$ is a control limit. In the remaining part of the paper, this CUSUM chart is called CUSUM-D chart, where "D" represents "drift".

From the construction of the CUSUM-D chart described above, it can be seen that the chart is constructed based on the log likelihood ratio statistic under the flexible OC framework that a time-varying shift of size $\delta_n \sigma_0$ occurs at time n. This OC framework contains all drift patterns, including the linear, quadratic and other types of drifts as special cases. To estimate δ_n by (1), the exponential weights $(1-\lambda)^{n-i}$ used in (1) are meaningfully non-zero only at a few time points close to n. Thus, the estimate $\widehat{\delta}_n$ is actually *local* to n. For the reasons explained above, the CUSUM-D chart should be robust to the mean drift pattern, which will be confirmed by a simulation example in the next section.

In the CUSUM-D chart (2)-(3), there are two parameters λ and h to determine in advance. As in a conventional CUSUM or EWMA chart, the value of λ can be specified in advance, and then the control limit h can be determined easily by simulation to achieve a given ARL_0 level. Selection of λ will be studied numerically in the next section. To evaluate the performance of the CUSUM-D chart, we can use the IC average run length, denoted as ARL_0, and the OC average run length, denoted as ARL_1, as discussed in Qiu (2014).

3 Simulation Study

The simulation study is organized in two parts. First, selection of the parameter λ used in (1) is studied carefully in a special case for detecting linear drifts in Subsection 3.1. Then, the IC and OC performance of the chart is studied in Subsection 3.2, in comparison with several representative existing methods.

3.1 Selection of λ used in (1).

We study the impact of λ on the performance of the CUSUM-D chart when detecting linear drifts in this part. Results in such cases can provide some guidelines for selecting λ in other cases. First, it should be pointed out that selection of λ may depend on the shift time τ and the drift slope γ. Consider cases when $\text{ARL}_0 = 200$, the IC process distribution is $N(0, 1)$, and the OC process distribution is $N(\gamma(n - \tau), 1)$, for $n > \tau$. In the above setup, γ is allowed to change its value among $\{0.01 * i, i = 1, \cdots, 30\}$, and τ is allowed to change among $\{20, 30, 40, 50, 60, 70\}$. For each combination of the γ and τ values, we search for the λ value so that the CUSUM-D chart has the smallest ARL_1 value, and the corresponding optimal λ value is denoted as λ_{opt}. In each case, the ARL_1 value is computed based on 10,000 replicated simulations. The results of the searched λ_{opt} values are shown in Figure 1. In the figure, the solid curve denotes $\lambda_{opt} = 2\gamma - 3\gamma^2$, which is obtained by the least square estimation when fitting a quadratic curve to the results shown in the last plot when $\tau = 70$. From the figure, we can have the following conclusions. First, when τ increases, the results will stabilize, and can be described well by the function $\lambda_{opt} = 2\gamma - 3\gamma^2$. This confirms that the CUSUM-D chart will enter the steady-state when the current observation time n increases, and the relationship between λ_{opt} and γ can be described well by the quadratic function $\lambda_{opt} = 2\gamma - 3\gamma^2$. Therefore, λ should be chosen larger when γ is larger, which is intuitively reasonable because the exponentially weighted least square regression procedure (1) should assign less weights to previous observations whose observation times are far away from the current time n for detecting a steeper mean drift. Second, when τ is small (e.g., $\tau = 40$), λ should be chosen according to the relationship $\lambda_{opt} = 2\gamma - 3\gamma^2$ for detecting relatively steep mean drifts (e.g., $\gamma > 0.1$), and 0 for detecting mean drifts that are quite flat, which is also intuitively reasonable.

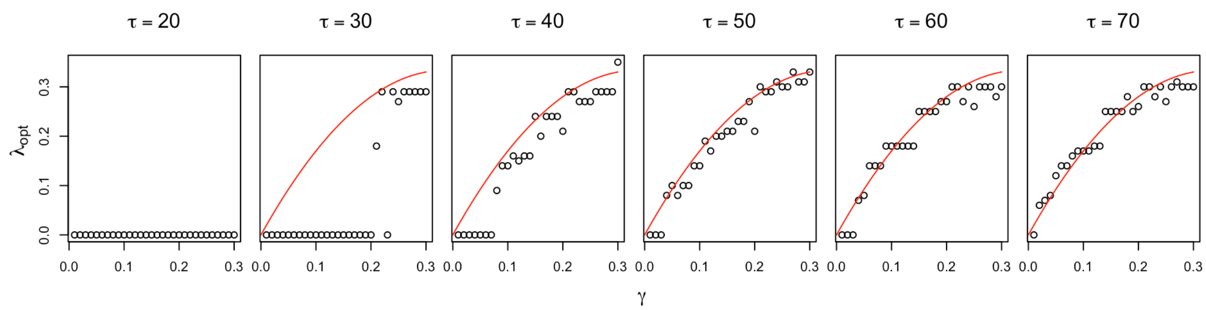

Figure 1: In each plot, small circles denote the optimal values of λ (i.e., λ_{opt}), and the solid quadratic line denotes $\lambda_{opt} = 2\gamma - 3\gamma^2$.

3.2 Comparison with several existing methods

In this part, we evaluate the numerical performance of our proposed control chart CUSUM-D, in comparison with several representative existing charts. The first existing chart is the conventional CUSUM chart, denoted as CUSUM (Page 1954). Readers are reminded that this conventional CUSUM chart is designed for detecting step mean shifts, and optimal under the assumptions that process observations at different time points are independent and identically normally distributed. The second existing chart is the GLR-L chart discussed in Zou et al. (2011), designed for detecting a linear mean drift using the classical likelihood ratio formulation for detecting a change-point. The third existing control chart considered is the ELM control chart proposed by Zhou et al. (2010).

This chart integrates the EWMA charting statistic with the generalized likelihood ratio statistic for monitoring a process with patterned mean or variance changes.

IC Performance. We first investigate the IC performance of the four charts. In the conventional CUSUM chart, the allowance constant k is chosen to be the commonly used value of 0.25. The weighting parameter is chosen to be 0.2 for the ELM chart and the CUSUM-D chart. The GLR-L chart is based on change-point detection, and does not have any parameter to choose. The assumed ARL_0 value is chosen to be 200 for all four charts and their actual ARL_0 values are computed by simulation as follows. First, we generate an IC dataset of size M from the $N(0,1)$ distribution. Then, the control limit is searched based on 5,000 bootstrap samples of the IC dataset until the assumed ARL_0 value (i.e., 200) is reached. Then, the chart is applied to the sequential observations generated from the IC model, and an IC run length value can be computed. This sequential monitoring process is repeated for 5,000 times, and the average of the 5,000 IC run length values is calculated as the actual ARL_0 value of the chart. We repeat the above process for 100 times. The average of the 100 ARL_0 values is used as the computed actual ARL_0 value of the chart. We consider cases when the IC sample size $M = 200$, 500, $1,000$, $2,000$ and $5,000$. The computed actual ARL_0 values of the four charts and their standard errors are presented in Table 1. From the table, it can be seen that i) the charts ELM and CUSUM-D have a reliable IC performance in all cases considered, ii) the computed actual ARL_0 value of the chart CUSUM is beyond 5% of the nominal ARL_0 value of 200 when $M = 200$ and within 5% of the nominal ARL_0 value when $M \geq 500$, and iii) the chart GLR-L has an unreliable IC performance when $M = 200$ and becomes reliable when $M \geq 500$. This example confirms that all four charts have a reliable IC performance when $M \geq 500$, and the charts ELM and CUSUM-D have a reliable IC performance even when $M = 200$.

OC Performance. To evaluate the OC performance of the four charts described above, we consider cases when $\tau = 50$, and $X_n \sim N(\gamma_1(n-\tau), 1)$ (i.e., linear mean drift) or $X_n \sim N(\gamma_2(n-\tau)^2, 1)$ (i.e., quadratic mean drift), for $n \geq \tau$, after the process becomes OC. Other setups

Table 1: Actual ARL_0 values and their standard errors (in parentheses) of the four charts CUSUM, GLR-L, ELM and CUSUM-D in cases when the IC sample size M = 200, 500, $1,000$, $2,000$, or $5,000$.

M	CUSUM	GLR-L	ELM	CUSUM-D
200	216.36(11.73)	128.27(6.55)	203.22(12.21)	205.87(11.51)
500	205.40(1.03)	191.67(9.41)	200.23(6.81)	201.48(6.95)
1,000	206.23(5.89)	197.74(5.96)	204.13(5.66)	203.48(5.65)
2,000	201.85(3.74)	197.49(3.77)	200.53(3.85)	200.58(3.93)
5,000	203.33(2.34)	199.91(2.59)	202.58(2.38)	202.12(2.42)

are the same as those in the example of Figure 1 above. For each chart, its nominal ARL_0 value is fixed at 200, and its control limit is chosen based on 10,000 IC process observations generated from the $N(0,1)$ distribution. To make the comparison among different charts fair, for detecting a given mean drift, the procedure parameters of each chart are chosen such that its ARL_1 value reaches the minimum. Namely, the optimal performance of the charts is compared here, which is recommended in the literature (cf., Qiu et al. 2020). In each case, the optimal ARL_1 value of a chart is computed based on 10,000 replicated simulations. The computed ARL_1 values of the four charts for detecting some linear mean drifts are presented in Table 2, and the corresponding results for detecting some quadratic mean drifts are presented in Table 3. From the tables, we can have the following conclusions. First, the proposed new chart CUSUM-D performs the best in all cases considered for detecting linear mean drifts, compared to the three competing methods. It is better than the conventional CUSUM chart because the latter is optimal for detecting step shifts only, and the former is designed for detecting mean drifts that is considered in this example. Second, it can be seen that the chart CUSUM-D is also the best among all four charts for detecting quadratic mean drifts. Thus, its performance is indeed quite robust to the mean drift patterns. Third, by checking the results of the other three charts, it can be seen that the charts GLR-L and ELM do not perform well, especially when the mean drifts are relatively small. This conclusion is actually consistent with the results in Zhou et al. (2010) and Zou et al. (2011) which confirmed that the charts GLR-L and ELM performed well only when the mean drifts were large. This example shows that CUSUM-D is indeed effective for detecting mean drifts and quite robust to the mean drift patterns.

Table 2: Optimal ARL_1 values with standard errors (in parenthesis) of the four control charts for detecting some linear mean drifts when $\tau = 50$ and $\text{ARL}_0 = 200$. Number in bold in each row denotes the smallest ARL_1 value in that row.

γ_1	CUSUM	GLR-L	ELM	CUSUM-D
0.001	98.93(0.72)	138.39(1.41)	104.79(0.99)	**81.65**(3.63)
0.005	52.30(0.38)	70.05(0.66)	53.61(0.48)	**43.57**(2.34)
0.01	37.65(0.31)	48.54(0.48)	38.11(0.35)	**32.74**(2.00)
0.05	16.11(0.22)	18.92(0.27)	16.22(0.20)	**15.38**(1.48)
0.1	10.87(0.21)	12.39(0.23)	10.92(0.18)	**10.67**(0.36)

Table 3: Optimal ARL_1 values with standard errors (in parenthesis) of the four control charts for detecting some quadratic mean drifts when $\tau = 50$ and $\text{ARL}_0 = 200$. Number in bold in each row denotes the smallest ARL_1 value in that row.

γ_2	CUSUM	GLR-L	ELM	CUSUM-D
0.0001	54.26(0.40)	63.90(0.58)	54.69(0.46)	**45.73**(2.39)
0.0005	32.82(0.30)	37.38(0.38)	32.89(0.30)	**29.48**(1.89)
0.001	26.06(0.26)	29.28(0.33)	26.12(0.26)	**24.03**(1.73)
0.005	14.88(0.23)	16.10(0.25)	14.90(0.20)	**14.44**(1.45)
0.01	11.61(0.21)	12.42(0.23)	11.65(0.19)	**11.44**(0.35)

4 An Application

In this section, we apply the proposed control chart CUSUM-D and the three alternative charts CUSUM, GLR-L and ELM to a real-data example about the daily exchange rates between Indian Rupee and US Dollar between Nov 01, 2010 and December 30, 2011. There are a total of 292 exchange rates during this time period, and they are shown in the left panel of Figure 2. From the plot, it can be seen that the first 200 observations are quite stable, compared with the last 92 observations, where the two groups of observations are separated by the vertical dashed line. So, the first 200 observations are used as IC data, and the remaining ones are used for online process monitoring.

From Section 2, the current version of the proposed CUSUM-D chart is constructed based on the assumptions that the process distribution is normal and the process observations at different time points are independent of each other. To check these assumptions, the R functions **shapiro.test()** and **Box.test()** are used for testing for the normality and autocorrelation of the IC data, respectively. The p-values of these two tests are 0.039 and < 0.001. Thus, the distribution of the IC data is significantly different from normal, and there is a significant serial data correlation in the IC data. To decorrelate the data, the serial data correlation structure is first estimated from the IC data, and then the recursive data decorrelation procedure discussed in Qiu et al. (2020) is used to decorrelate all observed data. To transform the decorrelated data so that the distribution of the transformed data becomes normal, we consider using the Johnson's transformation families (Johnson 1949, Slifker and Shapiro 1980) via the R function **jtrans()**. After data decorrelation and normalization, the R functions **shapiro.test()** and **Box.test()** are applied to the IC data again, and the p-values of these two tests are 0.936 and 0.887, respectively. Thus, the serial data correlation has been mostly deleted by the data decorrelation procedure, and the normality assumption is valid now after using the Johnson's transformation. The decorrelated and normalized data are shown in the right panel of Figure 2.

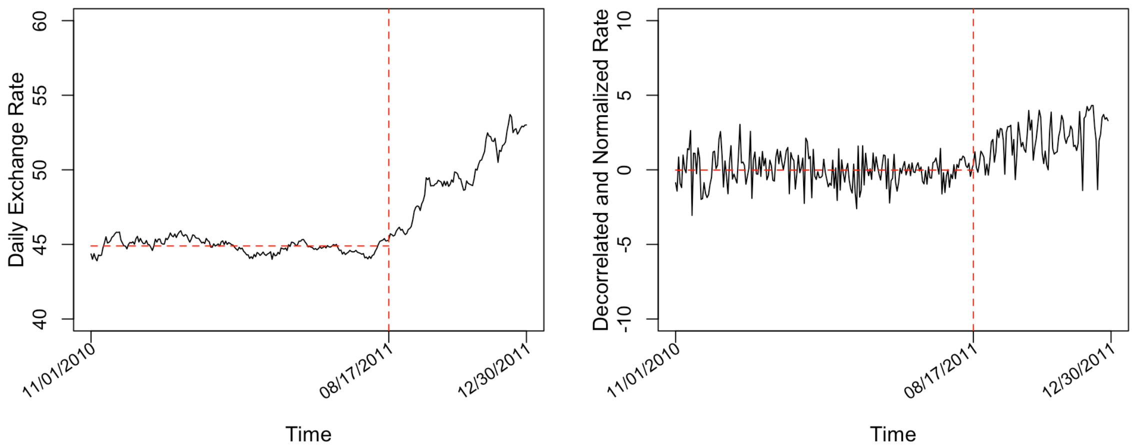

Figure 2: Left panel shows the original observations of the exchange rates between Indian Rupee and US Dollar during Nov 01, 2010 and December 30, 2011. Right panel shows the decorrelated and normalized data. The vertical dashed line in each plot separates the IC data and the data for online process monitoring, and the horizontal dashed line denotes the IC mean.

Next, we apply the proposed CUSUM-D chart to this data, together with the alternative charts CUSUM, GLR-L and ELM. In CUSUM-D, we set $\lambda = 0$ based on the simulation results shown in Figure 1 and the expectation that a mean drift can occur anytime soon after the online process monitoring starts. In the conventional CUSUM chart, the allowance constant k is chosen to be the commonly used value of 0.25. For the ELM control chart, the weighting parameter λ_e is chosen to be 0.2, as suggested in Zhou et al. (2010). The related charts are shown in Figure 3. The CUSUM, GLR-L, ELM and CUSUM-D charts give their first signals on 09/07/2011, 09/09/2011, 09/08/2011 and 09/06/2011, respectively. So, the CUSUM-D chart gives the earliest signal in this example. The detected upward drift should be related to the S&P downgrade of the credit rating of the United States from AAA to AA+ on August 5, 2011. This example shows that the CUSUM-D chart is quite effective in detecting such mean drifts.

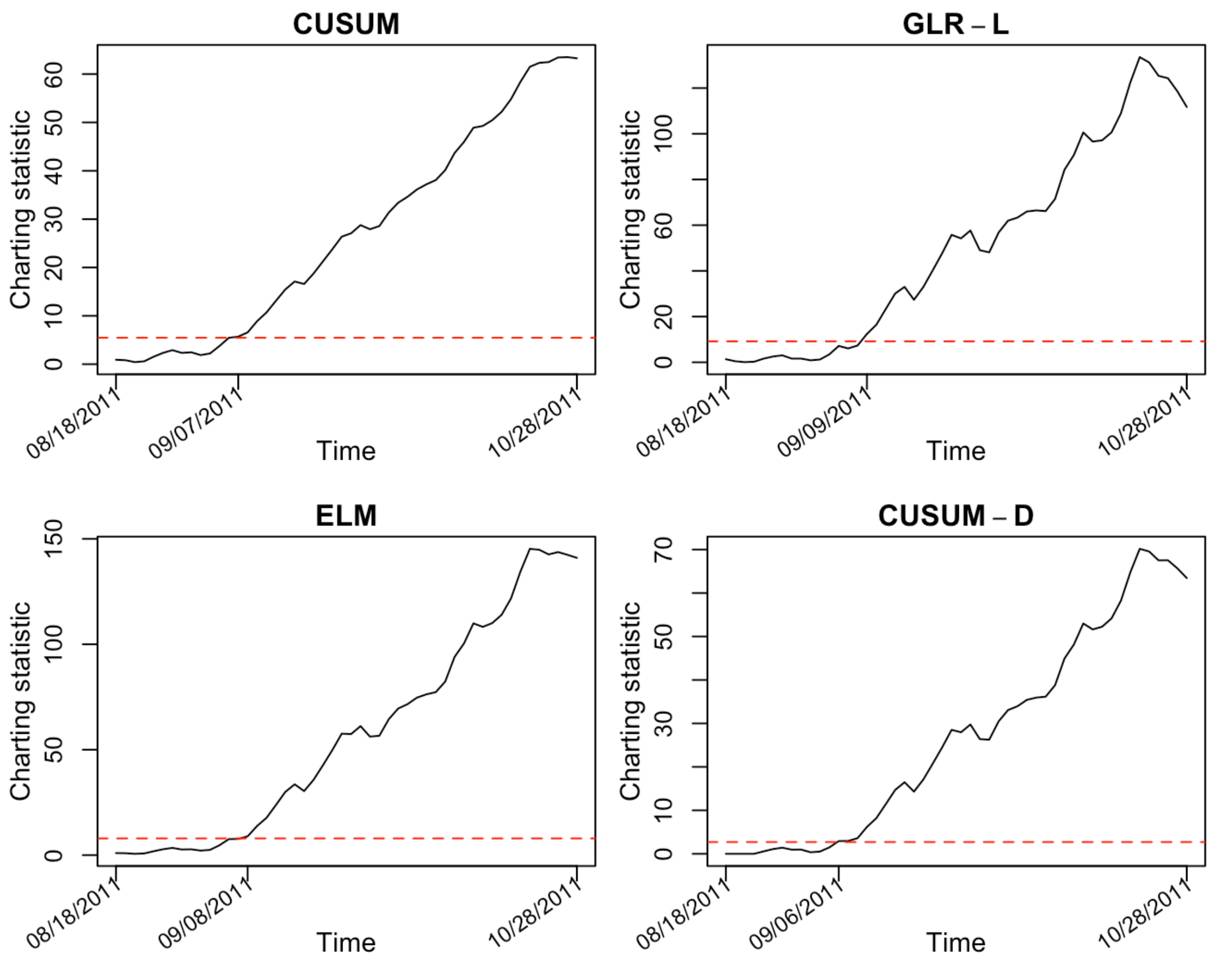

Figure 3: The four control charts when they are used for monitoring the decorrelated and normalized exchange rates between the Indian Rupee and US Dollar during 08/18/2011 and 10/28/2011.

5 Concluding Remarks

We have described a new control chart for detecting process mean drifts. The new method is constructed based on the generalized likelihood ratio statistic and the exponentially weighted least square estimation of a mean shift size. Numerical studies have confirmed that it is effective in detecting mean drifts in different cases considered. While our proposed chart CUSUM-D has been shown to perform the best, compared to its three peers, in the numerical studies presented in Section 3, we would like to remind the readers that this chart is designed for detecting mean drifts with unknown drift patterns. For detecting mean shifts, for example, the conventional CUSUM chart or its modified versions should be considered. The current version of the proposed method is designed for cases when the process distribution is normal and process observations at different time points are independent of each other. In cases when the normality assumption is invalid, a nonparametric version of the proposed method might be possible by using data categorization or other approaches for constructing nonparametric control charts (e.g., Qiu and Li 2011). In cases when the observed data are serially correlated, a time series model or a data decorrelation procedure by moment estimation of the serial data correlation might be used to remove the serial data correlation (e.g., Apley and Tsung 2002, Qiu et al. 2020) first, and then the proposed method can be used afterwards. Also, the steady-state relationship between λ_{opt} and γ for detecting linear drifts is obtained in Section 3.1 by numerical studies only. Theoretical justification of that result is needed. All these problems will be studied carefully in our future research.

Acknowledgments: The authors thank the editor and a referee for some constructive comments and suggestions which improved the quality of the paper greatly. This research is supported in part by an NSF grant with the grant number DMS-1914639.

References

Apley D.W., and Tsung, F. (2002), "The autoregressive T^2 chart for monitoring univariate autocorrelated processes," *Journal of Quality Technology*, **34**, 80–96.

Gan, F.F. (1992), "CUSUM control charts under linear drift," *The Statistician*, **41**, 71–84.

Han, D., and Tsung, F. (2004), "A generalized EWMA control chart and its comparison with the optimal EWMA, CUSUM and GLR schemes," *The Annals of Statistics*, **32**, 316–339.

Johnson, N.L. (1949), "Systems of frequency curves generated by methods of translation," *Biometrika*, **36**, 149–176.

Page, E.S. (1954), "Continuous inspection schemes," *Biometrika*, **41**, 100–115.

Qiu, P. (2014), *Introduction to Statistical Process Control*, Boca Raton, FL: Chapman Hall/CRC.

Qiu, P., and Li, Z. (2011), "On nonparametric statistical process control of univatiate processes," *Technometrics*, **53**, 390–405.

Qiu, P., Li, W., and Li, J. (2020), "A new process control chart for monitoring short-range serially correlated data," *Technometrics*, **62**, 71–83.

Shu, L., Jiang, W., and Tsui, K.L. (2008), "A weighted CUSUM chart for detecting patterned mean shifts," *Journal of Quality Technology*, **40**, 194–213.

Shuford, W.D., Warnock, N., Molina, K.C., and Sturm, K. (2002), "The Salton Sea as critical habitat to migratory and resident waterbirds," *Hydrobiologia*, **473**, 255–274.

Slifker J.F., and Shapiro, S.S. (1980), "The johnson system: selection and parameter estimation," *Technometrics*, **22**, 239–246.

Su, Y., Shu, L., and Tsui, K.L. (2011), "Adaptive EWMA procedures procedures for monitoring processes subject to linear drifts," *Computational Statistics and Data Analysis,* **55**, 2819–2829.

Tiffany, M.A., Ustin, S.L., and Hurlbert, S.H. (2007), "Sulfide irruptions and gypsum blooms in the Salton Sea as detected by satellite imagery, 1979-2006," *Lake and Reservoir Management*, **23**, 637–652.

Zeileis, A., Shah, A., and Patnaik, I. (2010), "Testing, monitoring, and dating structural changes in exchange rate regimes," *Computational Statistics & Data Analysis*, **54**, 1696–1706.

Zhou, Q., Luo, Y.Z., and Wang, Z. (2010), "A control chart based on likelihood ratio test for detecting patterned mean and variance shifts," *Computational Statistics and Data Analysis*, **54**, 1634–1645.

Zou, C., Liu, Y., and Wang, Z. (2009), "Comparisons of control schemes for monitoring the mean of processes subject to drifts," *Metrika*, **70**, 141–163.